

WHAT IS CLAIMED IS:

1. A method for servicing transmit traffic in a node of a network, the network including a plurality of nodes connected by first and second rings formed by two or more transmission media, the method comprising:

receiving usage data from a downstream node, the usage data including transit delay data associated with a plurality of downstream nodes;

receiving a packet for routing to the network;

determining a shortest path to a destination node including identifying one of the first and second rings as being associated with the shortest path;

determining if the identified one of the first and second rings is more congested than the other of the first and second rings using the transit delay data; and

if so, routing the packet to the destination on the other ring irrespective of the shortest path determination.

2. The method of claim 1, further comprising

determining transit delay data for the node;

appending the transit delay data for the node to the received transit delay data; and

forwarding the transit delay data including appended transit delay data to an upstream node.

3. The method of claim 2, wherein the step of appending transit delay data includes:

identifying transit delay data associated with a node farthest away from the node; and

dropping the transit delay data associated with the node farthest away from the node prior to appending the node's transit delay data.

4. The method of claim 1, wherein the transit delay data received is of the form of a plurality of vectors each reflecting the transit delay for their respective node.

5. The method of claim 4, wherein the step of receiving usage data includes receiving transit delay data from 32 downstream nodes.

6. The method of claim 1, further comprising determining a latency metric, the latency metric indicative of a delay between the node and the destination node, and the step of determining if a ring is more congested includes comparing the latency metrics associated with the destination node for each ring.

7. The method of claim 1, further comprising determining an average transit delay for each the plurality of nodes, the average transit delay computed as the average of a previously determined average transit delay for a given node and newly received delay data associated with the given node.

8. The method of claim 6 wherein the latency metric is computed as the mathematical average of a previously calculated latency metric indicative of a delay for nodes between the node and the given one of the plurality of downstream nodes and a newly calculated latency metric for a same path based on the received transit delay data.

9. The method of claim 7 wherein the step of determining if the identified one of the first and second rings is more congested than the other of the first and second rings includes using the average transit delay data computed for each of the plurality of downstream nodes.

10. The method of claim 1 further comprising
 recognizing when a packet may be part of a flow;
 storing flow information for a flow when a routing decision is made that routes a packet in a direction that is not consistent with the shortest path, the flow information including a flow direction selected and a timer;
 receiving another packet that is part of the flow;
 determining if a timeout period has expired since a last packet in the flow was sent based on the timer;
 if the timeout period has not expired , then routing the another packet to the destination based on the flow information including in a direction determined by the flow direction; and
 updating the timer to reflect a start of a new timeout period.

11. The method of claim 10 further comprising setting the timer to an initial value that is the greater of the latency period between the node and the destination node on both rings.

12. The method of claim 11 wherein the step of updating the timer includes setting the timer to a new value that is the greater of a current latency period between the node and the destination node on both rings.

13. The method of claim 1 further comprising determining if the destination node is farther away from the node than a predefined number of hops, and if so, routing the packet to the destination node based on the shortest path.

14. The method of claim 13 wherein the predefined number of hops is 32.

15. The method of claim 13 wherein a check is made to determine if a break has been detected in the network on one of the first and second rings, and if so, routing the packet to the destination node based on the shortest path.

16. The method of claim 1 further comprising storing in a table of destination nodes a hop count reflecting a hop count between the node and the given destination node for each of the first and second rings, a latency metric reflecting the congestion between the node and the given destination node for each of the first and second ring, a static ring selection based on the hop count, and a dynamic ring selection based on the latency metrics reflective of the congestion in the first and second rings between the node and the given destination node.

17. The method of claim 6 further comprising calculating the latency metric as the mathematical average of a previously calculated latency metric and an average transit delay for all nodes between the node and the given destination node.

18. The method of claim 17 wherein the average transit delay is weighted based on the number of hops between the node and the given destination node.

19. The method of claim 1 wherein the transit delay data is a measure of the amount of traffic in a low priority queue of a given downstream node.

20. A method for servicing transmit traffic in a node of a network, the network including a plurality of nodes connected by first and second rings formed by two or more transmission media, the method comprising:

receiving transit delay data associated with a plurality of downstream nodes;

receiving a packet for routing to a destination node in the network; and

routing the packet to one of the first and second rings not only based on a shortest path to the destination node but also on a transmit latency for both the first and second rings in view of the received transit delay data.

21. A node in a network including a plurality of nodes connected by first and second rings formed by two or more transmission media, the node comprising:

fairness logic configured to

receive usage data from a downstream node including transit delay data associated with a plurality of downstream nodes;

receive a packet from a host associated with the node for routing to the network;

determine a shortest path to a destination node including identifying one of the first and second rings as being associated with the shortest path;

determine if the identified one of the first and second rings is more congested than the other of the first and second rings using the transit delay data; and

if so, routing the packet to the destination on the other ring irrespective of the shortest path determination.

22. The node of claim 21 further comprising fairness logic configured to track flows associated with a node including remembering a last ring on which packets of the flow were forwarded to the node and setting a timer to a value reflective of a longest amount of time a packet will take to reach the node on either ring, receive a packet that is part of a flow and route the packet to the node using the last ring if the timer is unexpired.